

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 479 660 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention
of the grant of the patent:
23.10.1996 Bulletin 1996/43

(51) Int Cl.⁶: **G06F 9/46**

(21) Application number: **91402602.6**

(22) Date of filing: **30.09.1991**

(54) **Distributed configuration profile for computing system**

Verteiltes Konfigurationsprofil für ein Rechnersystem

Profil de configuration distribué pour système d'ordinateur

(84) Designated Contracting States:
DE FR GB IT NL

(30) Priority: **02.10.1990 US 591784**

(43) Date of publication of application:
08.04.1992 Bulletin 1992/15

(73) Proprietor: **DIGITAL EQUIPMENT CORPORATION**
Maynard, MA 01754 (US)

(72) Inventors:
• **Miller, Steven P.**
Newton, Massachusetts 02159 (US)

• **Butler, W. Lampson**
Cambridge, Massachusetts 02138 (US)

(74) Representative: **Dubois-Chabert, Guy et al**
Société de Protection des Inventions
25, rue de Ponthleu
75008 Paris (FR)

(56) References cited:
EP-A- 0 384 339

• **EXE vol. 2, no. 6, November 1987, U.K. pages 58**
- 61; Johnson, H.; Adams, M: 'RPC: the key to
distributed software.'

EP 0 479 660 B1

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

Description

This invention relates to distributed computing systems using data communications networks, and more particularly to methods of providing configuration information to clients in a distributed computing system regarding network services or remote procedure calls.

In U.S. Patent 4,823,122, issued April 28, 1989, assigned to Digital Equipment Corporation, assignee of this invention, a communications network is disclosed which is capable of interconnecting local computer terminals with remote service providers using a communications link. The local terminals are coupled to the communications link by interface units which handle network protocol functions for the terminals. One of the functions necessary at the local units is that of obtaining actual network addresses for services known locally only by service names. To perform this translation, service advertisements are periodically sent by the service providers, and a service directory is generated locally from these advertisements. Then, when a service request is generated locally, the service directory is searched to find the corresponding network address of the provider of this service. If there are multiple providers of the same service some mechanism must be provided for selecting the most efficient service provider for a local requester, and for distributing the load among the various providers in an efficient manner. In EP-A-0 384 339 a method is disclosed for allocating resources within a computer network by employing a broker mechanism. This broker mechanism operates by monitoring a subset of all available servers capable of delivering the requested service, and allocating based upon a network policy and available resources, to suggest the name of a provider to a service requester.

There are often redundant resources in a distributed system, in order to improve availability and share load. One problem faced by a user (i.e., an applications programmer) in such an environment is selecting a particular resource for use. The user would ideally like means of selecting an instance of a redundant resource that preserves the availability and load sharing, yet may be optimized to the particular environment and/or tailored to the user preferences. For example, when using an remote procedure call service (the resource) that supplies an employee address list, the remote procedure call user may prefer resources within the same LAN, with high speed communications available, to be selected in preference to resources available only via a wide area network. However, if the resources in the local LAN are not available, the remote procedure call user would like to use the remote resources accessible only via the wide area network. Furthermore, the designation of local and remote is relative. Each user, possibly in a different location, may have different views of what is local and what is remote. Therefore, the selection of resources must be able to be tailored to the user and or conditions local to the user.

An objective, then, in a system having redundant distributed resources, is to provide a mechanism that supports the customized selection of resources in a distributed environment, while maintaining availability and load sharing.

In accordance with one embodiment of the invention, a distributed computing system using a data communications network may have a number of service providers for a given service or remote procedure call. A client on the network makes reference to a name service to obtain the network address of one of these service providers. The name service maintains a configuration profile of the service providers in order to resolve the issue of selecting one of the several service providers when a request is made. A single configuration profile is a priority-ordered search list that maps from a service identifier (e.g., remote procedure call interface specification) into service provider (e.g., remote procedure call server) names. A configuration profile may include names for individual service providers, and/or named groups of service providers, and/or other configuration profiles. Configuration profiles are stored in a manner that makes them accessible throughout the distributed system, e.g. in the name service. Configuration profiles may be chained together by referencing other configuration profiles to provide a hierarchy of configuration profiles.

Configuration profiles address a number of related problems in a distributed system. This is regardless of the fact that the underlying name service is often structured as a single-keyed hierarchy, where the keys do not reflect the service. They provide mechanisms to automatically: (1) efficiently map a (flat) name of a service into a set of service providers; (2) customize the mapping of services into service providers based on a user, system, LAN, site, organization, etc.; (3) provide a hierarchical search path of these mappings; (4) support enhanced availability of service providers; (5) share load amongst service providers; (6) prioritize the selection of a provider within a profile; (7) randomize equal priority selections within a profile; (8) store the profiles in a manner available throughout a distributed system; (9) minimize the configuration management required for systems and users; and (10) minimize the use of well-known names. These mechanisms support user or site defined service utilization policies.

The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself, however, as well as other features and advantages thereof, will be best understood by reference to the detailed description of specific embodiments which follows, when read in conjunction with the accompanying drawings, wherein:

Figure 1 is an electrical diagram in block form of a distributed computing system employing a data communication network which may use features of the invention;

Figure 2 is a schematic representation of processors such as those in Figure 1 when executing a remote procedure call;

Figure 3 is a logical representation of operation of a remote procedure call as in Figure 3;

Figure 4 is a logical representation of a name service in the system of Figure 1, according to the invention;

Figure 5 is a logical representation of a configuration profile used with the name service of Figure 4, according to the invention;

Figure 6 is a logical representation of a service group entry in the name service of Figure 4; and

Figures 7 and 7a are logical representations of a server entry in the name service of Figure 4.

Referring to Figure 1, a computer network or distributed system is shown which may use features of the invention. A network including a communications link 10 is used to provide a communication path between a number of processors or computers including clients 11, each of which may be a single host computer or may be connected to a number of terminals or user nodes 12. The communications link may be a single local loop or bus (i.e., a local area network), or may be coupled by bridges 13 to more remote networks 10a, and/or by a microwave or satellite link 14 to other remote networks 10b (i.e., in a wide area network). Each one of the clients 11 (and the terminals 12) is assumed to be a processor executing its own code with its own local memory, and each processor may communicate with each of the others via the path 10, using a communications protocol which ordinarily involves message packets. Services are available on this network 10, 10a and 10b through servers 15 identified as server-1, server-2, server-3, etc. Each of the servers 15 provides one or more services 16 to the network, with these services being identified as service-A, service-B, service-C, etc. These services may be hardware facilities, such as high speed printers, back-up disks, or the like, or may be software facilities, such as name services, global data base utilities, or the like. As discussed below, the services 16 may be remote procedure calls. Note that the services 16 may be duplicated on the network. Service-A is seen to be available at both server-1 and server-3 on the local network 10, and at server-5 on the remote network 10a, as well as server-6 on the remote network 10b. A client 11 has access to service-A from any one of the servers 15 providing the service-A, but it may be more efficient to utilize the local server-1 or server-2 if not busy, since the response time is quicker and the communication facility is less burdened.

The clients 11 or user terminals 12 may be personal

computers, workstations, modems connected to remote users, video display terminals, point-of-sale terminals, automated teller machines, or any of a variety of such devices. The links 10, 10a or 10b between clients 11 and servers 15 may be local area networks using technologies such as Ethernet, token ring, FDDI, StarLan, or the like, or in some cases may be merely RS232 type or backplane connections. The physical construction of the communications links 10, 10a or 10b can be of various types, such as coaxial cable, twisted pair, fiber optics, or the like. The communication method and protocol employed can also be of various types, such as DECnet, OSI, or TCP/IP, for example (i.e., higher level protocol services). While these commonly-used communications methods are based on bit-serial transmission on the link 10, the features of the invention are applicable as well to communications methods using parallel data transmission on the link. Of course, the communications link 10 may include bulk transfer facilities (trunk lines) using multiplexing, such as T1 circuits or the like. Or, a bulk transfer facility may be one of the services represented by the services 16. The communications network represented in Figure 1 may service hundreds or thousands of clients, user terminals and servers, depending upon the installation. Since the network 10 of Figure 1 may be connected to other such networks 10a and 10b by bridges and/or wireless links (either represented by services 16 or otherwise, i.e., transparent to the protocol), the number of users and the number of service providers can be quite large.

In a preferred embodiment of the invention, each of the processors 11 and 12 and the servers 15 is able to make remote procedure calls to the other processors, and the services 16 may themselves be remote procedures which may be called from other processors on the network; alternatively, remote procedure calls executed on the servers 15 may be used to access the services 16 (as well as for other purposes). A remote procedure call (RPC) is a procedure call mechanism, extended to provide transfer of control and/or data across communication networks such as the network 10 of Figure 1. Referring to Figure 2, if a process A (or a thread within a process) is executing in processor 11a, it can make a local call to procedure B in the address space of processor 11a, or make a remote call to a procedure C in the address space of processor 15a via RPC facility 17 and network 10 using the communications layer 18. Execution of the process A is suspended when the remote call is made, and then the procedure C executes, returning data to the process A where execution recommences. Note that for remote procedure calls, there is no shared memory; the procedures are ordinarily in two physically separate computers. Therefore, the parameters passed back and forth must be actual values rather than pointers, i.e., the data must be copied from machine to machine. Ordinarily the RPC mechanism is transparent to the user (i.e., to the writer of source code for an applications program); the structure of the call is the same

(from the standpoint of the process A) whether the call is local or remote. When the call is actually executed, however, the system must introduce an address of the remote procedure C so that message packets can be made up by the communication layer to send over the network 10. Providing this address is called "binding".

Referring to Figure 3, a program structure for remote procedure calls is based upon the concept of stubs. When making a remote call, five pieces of program are involved: the user program A (as in Figure 2), the user stub A_g, and one instance of RPCRuntime 17 execute in the caller machine 11a. The server procedure C, the server stub C_g, and another instance of RPCRuntime 17 execute in the callee machine 15a. When the user program A wishes to make a remote procedure call, it makes a local call which invokes a corresponding procedure in the user stub A_g. The user stub is responsible for placing a specification of the target procedure and arguments into one or more packets and asking RPCRuntime 17 to transmit these reliably to the callee machine 15a via network 10. On receipt of these packets, the RPCRuntime 17 in the callee machine 15a passes them to the server stub C_g, which unpacks them and makes a local call to invoke the appropriate procedure in the server procedure C. Meanwhile, the calling process A in the caller machine is suspended awaiting a result packet. When the call in the server procedure C completes, it returns to the server stub C_g and the results are passed back to the suspended process A in the caller machine 11a, via the user stub A_g which unpacks them. The RPCRuntime program 17 is responsible for retransmissions, acknowledgements, packet routing, and encryption, either directly or in combination with lower level services offered in the network. Apart from the effects of intermachine binding, the call happens just as if the user had invoked the procedure in the server directly; indeed, if the user code A and server code C were brought into a single machine and bound directly together without the stubs A_g and C_g the program would still work. The RPCRuntime program is described by Birrell et al in "Implementing Remote Procedure Calls", ACM Transactions on Computer Systems, Feb. 1984, pp. 39-50. The remote interactions are defined in an interface declaration, which is essentially a set of (abstract) procedure definitions of the remote service or program, combined with other relevant specification information. It defines the remote interface between the user and provider of the remote service; the interface declaration should be expressed in an interface specification language which allows heterogeneous implementations of the programs which provide and use the declared interface. The programmer of the user code and server code need not be concerned with creating the stubs or with the communications of packing and unpacking, but merely must avoid specifying arguments or results that are incompatible with the lack of shared address space. The programmer must also take steps to invoke intermachine binding, and to resolve occurrence-

es of network failure or failure of the server.

There are two aspects of binding to consider. First, the client or caller machine 11a must specify what service it is to be bound to, and this is a question of naming. Secondly, the client must determine the machine address of a particular callee or server, which is a question of location, and specify to this callee the procedure to be invoked. The binding operation is to bind an importer of an interface to an exporter of an interface; after binding, calls made by the importer invoke procedures implemented by the remote exporter.

The communications network 10 and communications layers 18 illustrated in Figures 1 and 2, in one embodiment, might be generally of the type disclosed in the above-mentioned U.S. Patent 4,823,122. this type of network executes a local area protocol. The communications protocol described in the patent may make use of virtual circuits to define a two-way path between a particular server client and a particular node or server 15. A feature of the protocol described in the U.S. Patent 4,823,122 is the use of advertising messages sent periodically by the service providers to advise all clients of what services are available on the network. To this end, a service advertising message is sent periodically by each service provider, by which it identifies itself (by network address) and the service or services provided at its node. The service advertising message contains information such as identification of the transmitting node (one of the servers 15, for example), a type field identifying this message as a service advertising message, and a multi-cast address field enabling all of the nodes to receive the message if desired. In the body of the advertising message each of the one or more services provided at this node is identified. The method described in the patent, however, places a burden on the communications link and imposes overhead upon each client node in keeping track of the service providers, especially in wide area networks having a large number of stations.

To provide the actual network addresses of facilities on the network 10, a name service 20 is provided as one of the services 16 on a network of Figure 1, available to all users via a server 15. The network address of the name server 20 is always made available to a user 11, 12 or 15 upon sign-on or boot-up, so inquiries can be made to the name server 20 to obtain a network address for any service or procedure or function for which the user has a name or an alias or generic name. A primitive way of using the name service 20 would be for it to collect service advertisements as discussed above, rather than this burden being imposed upon local clients 11. Another way is for the service providers to export their availability, i.e., export an interface identifier as described below. In any event, the name service is a commonly-used utility on communications networks, and acts as a database, collecting actual network addresses for services (or other named elements) on the network and responding to inquiries by looking up data in its da-

tabase and returning information found.

In the network of Figure 1, where a number of instances of various services, such as service-A, may be available at various physical locations, the name service 20 or RPC client must have some mechanism for selecting one or the other of these instances when an inquiry is made. For example, when an application is distributed, copies of the server (services 16) are installed on several systems (servers 15) whose locations are unknown to clients 11. To use a service, a client 11 must find a compatible server 15, that is, a server that offers a given RPC interface and any resource, such as a database, that the client 11 requests. Services 16 (via their servers 15) use the name service 20 to advertise themselves to clients 11, this name service being a repository of information about computing resources. Referring to Figure 4, the name service 20 maintains a network-wide database 21 called a namespace which stores information for applications for each distributed computing environment. Note that there may be more than one distributed computing environment, and thus more than one name service 20 and namespace 21, existing on a network such as that of Figure 1 at a given time, but only one will be treated herein; also, the name service database may be partitioned or replicated - it need not be one server. To find information in the namespace 21, a client 11 must access the name service 20 that maintains the namespace 21. Information about a service 16 and its applications is placed in a server entry 22 in the namespace 21, by an operation invoked by system management, by an applications program, or by the server exporting the information. This information includes location (network address), and the RPC interfaces and resources it offers. The name server 20 handles requests from clients 11 (received via messages over link 10) for information, returning (via messages) the information about a service 16 that satisfies the client's needs, if such an entry 22 exists.

According to the invention, an entry maintained by the name service 20 for a client 11 (or group of clients), to list a subset of the services 16 available to this client, may be a "configuration profile" as seen in Figure 5. A configuration profile 23 is a data entry in the namespace 21 in the form of a priority-ordered search list, keyed by service so it can contain information for multiple services 16. Each item 23a, 23b, etc., in a configuration profile 23 represents a service 16 and may contain server entries 22, service group entries 24, or other configuration profile entries 25 (as described below).

When more than one instance of a service is available on a network, the naming service 20 may be set up by an applications manager to define a service group entry 24 as seen in Figures 5 and 6. A service group is composed of the identities of an arbitrary set of the servers 15 which offer a service 16; the applications manager decides that this set should be treated as a group. In addition, a service group entry 24 can include the identity of another service group entry 24a. Defining one

or more service groups for a service frees clients 11 from dependence on any one server 15. When a client 11 requests an RPC interface which corresponds to a service group, the name service 20 looks up (in the service group entry 24) information about the alternative servers, selects a compatible server 15 according to some predefined selection algorithm, and returns the binding (network, address) to the client 11. The algorithm is designed to share the load across a number of potential servers.

The name service 20 supports configuration profiles 23, according to the invention. Referring to Figure 5, a configuration profile 22 provides a mechanism for directing lookup operations in a name service. This lookup in a configuration profile is by service (i.e., RPC interface ID) requested, rather than by the name of server or object. A configuration profile 22 contains information about the names of servers 15 that offer a given RPC interface and associated resources. In a configuration profile 23, the ID of an RPC interface corresponds to a list of server entries 22, service group entries 24, and/or other configuration profile entries 25. In turn, the service group entries 24 contain server entries as described above for service groups, and the configuration profile entries contain server entries, service group entries and/or still other configuration profile entries.

A namespace configuration profile entry 23 in the name service 20 can thus be of three types (or a composite of these three types), as described: (1) a server entry 22 corresponding to a single instance of a server 15 running in its own address space - as part of a server's startup, it creates or verifies a server entry 22 in a name service 20, placing the ID and versions of each interface it offers in that entry, along with at least one binding for the interface, and optionally, a server entry 22 can also contain the object ID of any resource that the server can access; (2) a service-group entry 24 identifying a set of servers; (3) a configuration-profile entry identifying servers, service groups and/or other configuration profiles for an interface. When all types of entries are available, a client 11 importing a service interface searches for bindings by processing entries in the following order: (1) server entries, (2) service groups, and (3) configuration profiles. Although this order of processing is preferred, another order may be selected.

An object in the namespace 21 may contain many attributes, some created by RPC services, some created by other applications. Thus, while one attribute qualifies an object as a server entry, another may qualify the same object as a service-group entry, a configuration-profile entry, or some other non-RPC entry. An item 22 in the namespace 21 thus has at least one attribute field to qualify the object. One of these attribute fields may be a "priority" attribute, so that when the information is returned to the client the entries can be selected according to priority. For example, if there are several servers identified a priority 1, these will be tried first (in pseudorandom or perhaps some deterministic order), and if

none of these satisfy the requirements for supplying the service, then the ones having priority-2 attribute are tried, etc. The priority attributes are ordinarily added or modified by the manager, although the client can be given access to generate his own configuration profile with selected priority attributes.

A server entry 22 is illustrated in more detail in Figures 7 and 7a. A server entry may be merely one interface ID and its binding as seen in Figure 7, or may have multiple interface IDs and bindings as seen in Figure 7a. The name service 20 allows an application to create and reference each namespace entry 22 by name. For each RPC interface exported by an RPC server 15, a server entry 22 maintained by the name service 20 contains (1) one or more interface IDs, and (2) one or more bindings to the server 15, where a binding is a network address. A single entry 22 can contain multiple IDs and versions, each with its own set of bindings, and, in addition, a server entry optionally can contain a set of RPC object IDs, each of which corresponds to a specific resource offered by the server 15, as seen in Figure 7a. When importing any RPC interface offered by the server, a client 11 can request any of the object IDs in a server entry. To import an RPC interface a client specifies a name-service entry (or group, or service group, or interface).

The name service operations used in connection with the configuration profiles of the invention are: (1) exporting, (2) managing service groups and configuration profiles, (3) importing, and (4) unexporting.

Exporting is an operation by which a server 15 places information about an interface in a designated namespace 21, to publicly advertise a service 16 for use by any application running on a client 11. Optionally, the server can specify the object IDs of resources offered by that service. Unexporting is an operation by which a server 15 removes one or more bindings (or resources) from a server entry 22.

In managing service groups and configuration profiles, clients 11, servers 15 and independent management modules (not shown) can use the name service 20 for accessing a namespace 21 to create, update, look up, and remove server entries 22, service group entries 24 and configuration profile entries 23, subject to access control.

Importing is an operation by which the name service finds a binding that meets specific criteria specified by a client and returns that binding to the client. The selection is under some circumstances quasi-random, under others deterministic.

A feature needed in the management of configuration profiles is a mechanism for cycle checking, meaning that when a configuration profile is added or modified, the contents of each nested configuration profile are compared to make sure one doesn't contain an entry referring back to another configuration profile in the same chain, which would produce a condition of an infinite loop.

As thus described, a configuration profile is an optional, set-valued attribute of any entry in the distributed namespace 21. Each set member specifies a service or resource 16, via an interfaceUUID (Universal Unique Identifier) and interface version, acting as a lookup key within the profile. A comment field is used to contain the name of the interface. The name of the service interface is derived from the name attribute in an RPC interface definition file; this is a local name, possibly ambiguous. Each member also specifies the name of a service provider, and a priority. It is expected that profiles will most frequently be used for RPC services, although other services or resources can use them by simply creating their own UUIDs. Service provider names refer either to individual servers, service groups, or other profiles. Name service entries for individual servers contain their protocol towers (a data structure representing the address of the server and the hierarchy of network protocols used), including addresses; entries for service groups contain a list of 'equivalent' servers (which can contain names of other service groups); servers, service groups, and profiles are all represented in the name service as distinguished attributes on any name service entry. The priority value specifies a search priority that can be used to group members of equal priority. The nullUUID interface is used as a default profile pointer.

There are two categories of profile operations. The first are the primitive application program interface operations used to manipulate the profiles -- create the attributes, add members, delete members, read members, etc. These are all, conceptually, management operations done by a knowledgeable site or group or organization administrator. But they could be embedded in actual RPC servers.

The second category of profile operations are the query (import or lookup) operations that resolve a profile request. The query operations require a single piece of configuration information (provided at boot-up by boot code, or provided by a manual entry) which is the starting profile. First, the starting profile is read. If it contains members for the requested service, the highest priority members listed are grouped together. Service groups and other profiles are elaborated, in a specified order. From the equal priority grouping, a semi-random selection is returned. If the Import or other query operation is satisfied, the query is done. If not, subsequent selections from the equal priority group are returned. This process continues until either the query is satisfied or all equal priority members are exhausted. If the latter, then the next lower priority subset is gathered, and the selection iteration continues.

If a profile member points to another profile, it will be similarly elaborated. Profiles can be chained together in a hierarchy in this fashion. If none of the members for the specified interface satisfy the query, or there are no members for the specified interface, the query continues via following the default interface (null UUID) profile members. The query operation detects cycles; this can-

not simply be done when profiles are created, since the native name service access could also be used to manipulate profiles, but doesn't understand their semantics so wouldn't be able to check for cycles.

Configuration profiles are transparent to server applications and are transparent to client applications using RPC import or lookup operations. These client applications simply denote the service desired, and the operations begin the profile search with the starting profile. This -- the starting profile -- is the only one piece of information required for access to numerous services. It can be explicitly specified or implicitly read from an environment variable or similar mechanism.

Thus when a new system is unpacked and connected to the network, the user must learn and enter into his system a starting profile. This would typically be the user's or system's profile, and would be null except for a default profile entry pointing to a group or site profile already defined. The site profile would normally include a default organization profile, etc. So with this one piece of configuration information, the new system can make use of any of the RPC services available through the entire defined search tree hierarchy of profiles.

Since the configuration profiles are stored in the distributed name service, they greatly simplify the installation of new systems. A distributed system manager can install a new system in the environment and give it access to all the resources by simply providing the new system the name of its starting profile for the environment. Similarly, a new user can be configured by creating a new profile that is null except to link to a default profile as its parent.

Claims

1. A method of operating a distributed computing system, comprising the steps of:

connecting a number of clients and a number of service providers to a communications network, each service provider supplying a service having a service name, where each client is able to send data to and receive data from each of said service providers via said network;

maintaining a name service as one of said service providers, said name service containing at least one configuration profile associated with at least one of said clients, each configuration profile including a priority-ordered search list mapping service names to service providers;

defining for at least one of said clients a selected one of said configuration profiles for use by said client;

making a request from one of said clients to

said name service using one of said service names, and returning to said client information from said selected one of said configuration profiles for selecting from said search list one of said service providers and thereby providing to said one client a network address of said selected service provider;

making a request from said client to said selected service provider for said service.

2. A method according to claim 1 wherein said services include remote procedure calls and said service providers for remote procedure calls are mapped as interface specifications in said configuration profile.

3. A method according to claim 1 wherein each of said configuration profiles comprises a plurality of entries including:

optionally, one or more server entries each defining a network address of one of said service providers;

optionally, one or more service groups, each service group defining a plurality of said server entries; and

optionally, one or more other configuration profiles.

4. A method according to claim 1 wherein said service names include alphabetical names and numeric names.

5. A method according to claim 1 wherein said communications network includes a local area network and a wide area network, and wherein at least some of said services are duplicated on said local area network and/or on said wide area network.

6. A distributed computing system, comprising:

a communications network having a number of clients and a number of service providers connected thereto, each service provider supplying one or more services, wherein each client is able to send data to and receive data from each of said service providers via said network using a unique network address;

one of said service providers providing a name service, said name service being a database service containing entries for each of said service providers and receiving inquiries from said clients for the identity of said service providers referenced by service name and returning in-

formation identifying said service providers by network address;

said name service including in said database a configuration profile for at least one client or group of clients, each configuration profile comprising a plurality of entries including:

optionally, one or more server entries each defining a network address of one of said service providers;

optionally, one or more service groups, each service group defining a plurality of said server entries; and

optionally, one or more other configuration profiles;

each of said entries of each said configuration profile of said name service containing priority information to determine the order of use of said network addresses by said client after receipt of said inquiries by said name service;

and means for identifying for said client a selected one or more of said configuration profiles for reference upon making said inquiries.

7. A system according to claim 6 wherein said communications network includes a local area network and a wide area network, and wherein at least some of said services are duplicated on said local area network and/or on said wide area network.

8. A system according to claim 6 wherein said services include remote procedure calls, and said server entries include interface names for said remote procedure calls.

9. A system according to claim 6 wherein said entries of each said configuration profile make up a priority-ordered search list.

10. A name service provider for use in a distributed computing system, the system including a communications network having a number of clients and a number of service providers connected thereto, each of said service providers supplying a service identified by a service name, said name service provider comprising:

a database service containing entries for each of a plurality of services supplied by said service providers;

means for receiving inquiries from said clients for the identity of said service providers refer-

enced by service name and returning information identifying said service providers by network address;

said means including a configuration profile included in said database for one of said clients or for a group of said clients, each said configuration profile comprising a plurality of entries including:

optionally, one or more server entries each defining a network address of one of said service providers;

optionally, one or more service groups, each service group defining a plurality of said server entries; and

optionally, one or more other configuration profiles;

each of said entries of each said configuration profile of said name service containing priority information to determine the order of use of said network addresses by said client;

and means for identifying for said client a selected one or more of said configuration profiles for reference upon making said inquiries.

11. A name service provider according to claim 10 in combination with a communications network including a local area network and a wide area network, and wherein at least some of said services are duplicated on said local area network and/or on said wide area network.

12. A name service provider according to claim 10 wherein said services include remote procedure calls, and said server entries include interface names for said remote procedure calls.

13. A name service provider according to claim 10 wherein said entries of each said configuration profile make up a priority-ordered search list.

14. A method of operating a distributed computing system, comprising the steps of:

connecting a number of clients and a number of service providers to a communications network, each service provider supplying one or more services identified by a service name, wherein each client is able to send data to and receive data from each of said service providers via said network using a unique network address;

providing a name service as one of said service providers, and storing in said name service entries for each of said service providers;

receiving by said name service inquiries from said clients for the identity of said service providers referenced by service name and returning information from said name service to said clients identifying said service providers by network address in response to said inquiries;

said name service maintaining a configuration profile for at least one client or group of clients, each configuration profile storing a plurality of entries including:

optionally, one or more server entries each defining a network address of one of said service providers;

optionally, one or more service groups, each service group defining a plurality of said server entries; and

optionally, one or more other configuration profiles;

and storing in each of said entries of said configuration profile of said name service priority information to determine the order of use of said network addresses by said client.

15. A method according to claim 14 wherein said communications network includes a local area network and a wide area network, and wherein at least some of said services are duplicated on said local area network and/or on said wide area network.

16. A method according to claim 14 wherein said services includes remote procedure calls, and said server entries include interface names for said remote procedure calls.

17. A method according to claim 14 wherein said entries of said configuration profile make up a priority-ordered search list.

18. A communications network comprising:

a) a communications link;

b) a plurality of clients coupled to said communications link;

c) a plurality of service providers coupled to said communications network;

d) each one of said clients and each one of said

service providers including processing means executing a separate instruction stream, and each one of said clients and each one of said service providers having a unique network address on said network;

e) each one of said clients and each one of said service providers having means for making remote procedure calls to said service providers via said communications link, when executing said instruction stream;

f) each one of said clients and each one of said service providers having means for sending and receiving messages to other clients and service providers via said communications link using said network addresses, said messages being employed in making said remote procedure calls;

g) one of said service providers being a name service functioning to return a message containing the network address of a particular one of said service providers in reply to a message from one of said clients giving the name of said service;

h) and means included in said name service to maintain a configuration profile for at least one of said clients or group of clients, each configuration profile comprising a priority-ordered search list mapping service names into service providers, each configuration profile including:

optionally, one or more server entries each defining a network address of a service provider for a given service name;

optionally, one or more named service groups each defining a plurality of said server entries; and

optionally, one or more other said configuration profiles;

i) means for associating one of said configuration profiles with a given client;

j) and means in said name service to return said network addresses to said given client upon receipt of inquiries by said name service from said given client, for use by said client according to priority order of said list.

19. A network according to claim 18 wherein said service provider is identified in said service entry as a remote procedure call interface specification.

20. A network according to claim 18 wherein said communications network includes a local area network and a wide area network, and wherein at least some of said services are duplicated on said local area network and/or on said wide area network.

5

21. A network according to claim 18 wherein said service names include alphabetical names and numeric names.

10

Patentansprüche

1. Verfahren zum Betreiben eines verteilten Computersystems, mit den Schritten:

15

Anschließen einer Anzahl von Clients und einer Anzahl von Diensteanbietern an ein Kommunikationsnetz, wobei jeder Diensteanbieter einen Dienst liefert, der einen Dienstnamen besitzt, wobei jeder Client über das Netz an jeden Diensteanbieter Daten senden kann und von jedem Diensteanbieter Daten empfangen kann;

20

Bereithalten eines Namensdienstes als einen Dienst der Diensteanbieter, wobei der Namensdienst wenigstens ein Konfigurationsprofil enthält, das mit wenigstens einem der Clients in Beziehung steht, wobei jedes Konfigurationsprofil eine nach Priorität geordnete Suchliste enthält, die Dienstnamen auf Diensteanbieter abbildet;

25

Definieren eines ausgewählten der Konfigurationsprofile für wenigstens einen der Clients für die Verwendung durch diesen Client;

30

Tätigen einer Anforderung von einem der Clients an den Namensdienst unter Verwendung eines der Dienstnamen und Zurückschicken von Informationen von dem ausgewählten der Konfigurationsprofile zum Client, um aus der Suchliste einen der Diensteanbieter auszuwählen und dadurch für den einen Client eine Netzadresse des ausgewählten Diensteanbieters bereitzustellen;

40

45

Tätigen einer Anforderung dieses Dienstes von dem Client an den ausgewählten Diensteanbieter.

50

2. Verfahren nach Anspruch 1, in dem die Dienste Fernprozeduraufrufe enthalten und die Diensteanbieter für Fernprozeduraufrufe als Schnittstellenspezifikationen im Konfigurationsprofil abgebildet sind.

55

3. Verfahren nach Anspruch 1, in dem jedes der Kon-

figurationsprofile mehrere Einträge umfaßt, einschließlich:

optional eines oder mehrerer Server-Einträge, wovon jeder eine Netzadresse eines der Diensteanbieter definiert;

optional einer oder mehrerer Dienstgruppen, wobei jede Dienstgruppe mehrere der Server-Einträge definiert; und

optional eines oder mehrerer anderer Konfigurationsprofile.

4. Verfahren nach Anspruch 1, in dem die Dienstnamen alphabetische Namen und numerische Namen enthalten.

5. Verfahren nach Anspruch 1, in dem das Kommunikationsnetz ein lokales Netz und ein Weitverkehrsnetz enthält und in dem wenigstens einige der Dienste im lokalen Netz und/oder im Weitverkehrsnetz dupliziert werden.

6. Verteiltes Computersystem, mit:

einem Kommunikationsnetz mit einer Anzahl von Clients und einer Anzahl von Diensteanbietern, die daran angeschlossen sind, wobei jeder Diensteanbieter einen oder mehrere Dienste liefert, wobei jeder Client über das Netz unter Verwendung einer eindeutigen Netzadresse Daten an jeden der Diensteanbieter senden und Daten von jedem der Diensteanbieter empfangen kann;

wobei einer der Diensteanbieter einen Namensdienst anbietet, wobei der Namensdienst ein Datenbankdienst ist, der Einträge für jeden der Diensteanbieter enthält und Abfragen von den Clients hinsichtlich der Identität der Diensteanbieter, die mit dem Dienstnamen bezeichnet sind, empfängt und Informationen, die die Diensteanbieter durch die Netzadresse identifizieren, zurückschickt;

wobei der Namensdienst in der Datenbank ein Konfigurationsprofil für wenigstens einen Client oder eine Gruppe von Clients enthält, wobei jedes Konfigurationsprofil mehrere Einträge umfaßt, einschließlich:

optional eines oder mehrerer Server-Einträge, wovon jeder eine Netzadresse eines der Diensteanbieter definiert;

optional einer oder mehrerer Dienstgruppen, wobei jede Dienstgruppe mehrere

der Server-Einträge definiert; und

optional eines oder mehrerer anderer Konfigurationsprofile;

5

wobei jeder der Einträge jedes der Konfigurationsprofile des Namensdienstes Prioritätsinformationen enthält, um die Reihenfolge der Verwendung der Netzadressen durch den Client nach dem Empfang der Abfragen durch den Namensdienst zu bestimmen;

10

und eine Einrichtung zum Identifizieren eines oder mehrerer ausgewählter Konfigurationsprofile für den Client für die Bezugnahme nach dem Tätigen der Abfragen.

15

7. System nach Anspruch 6, in dem das Kommunikationsnetz ein lokales Netz und ein Weitverkehrsnetz enthält und in dem wenigstens einige der Dienste im lokalen Netz und/oder im Weitverkehrsnetz dupliziert werden.

20

8. System nach Anspruch 6, in dem die Dienste Fernprozeduraufrufe enthalten und die Server-Einträge Schnittstellennamen für die Fernprozeduraufrufe enthalten.

25

9. System nach Anspruch 6, in dem die Einträge jedes der Konfigurationsprofile eine nach Priorität geordnete Suchliste bilden.

30

10. Namensdiensteanbieter für die Verwendung in einem verteilten Computersystem, wobei das System ein Kommunikationsnetz mit einer Anzahl von Clients und einer Anzahl von Diensteanbietern, die daran angeschlossen sind, enthält, wobei jeder der Diensteanbieter einen Dienst liefert, der durch einen Dienstnamen identifiziert ist, wobei der Namensdiensteanbieter enthält:

35

40

einen Datenbankdienst, der Einträge für jeden der mehreren Dienste enthält, die von den Diensteanbietern geliefert werden;

45

eine Einrichtung zum Empfangen von Abfragen von den Clients hinsichtlich der Identität der Diensteanbieter, die mit dem Dienstnamen bezeichnet sind, sowie zum Zurückschicken von Informationen, die die Diensteanbieter durch die Netzadresse identifizieren;

50

wobei die Einrichtung ein in der Datenbank enthaltenes Konfigurationsprofil für einen der Clients oder für eine Gruppe der Clients enthält, wobei jedes Konfigurationsprofil mehrere Einträge umfaßt, einschließlich:

55

optional eines oder mehrerer Server-Einträge, wovon jeder eine Netzadresse eines der Diensteanbieter definiert;

optional einer oder mehrerer Dienstegruppen, wobei jede Dienstegruppe mehrere der Server-Einträge definiert; und

optional eines oder mehrerer anderer Konfigurationsprofile;

wobei jeder der Einträge jedes der Konfigurationsprofile des Namensdienstes Prioritätsinformationen enthält, um die Reihenfolge der Verwendung der Netzadressen durch den Client zu bestimmen;

und eine Einrichtung zum Identifizieren eines oder mehrerer ausgewählter der Konfigurationsprofile für den Client für die Bezugnahme beim Tätigen der Abfragen.

11. Namensdiensteanbieter nach Anspruch 10 in Kombination mit einem Kommunikationsnetz, das ein lokales Netz und ein Weitverkehrsnetz enthält, wobei wenigstens einige der Dienste im lokalen Netz und/oder im Weitverkehrsnetz dupliziert werden.

12. Namensdiensteanbieter nach Anspruch 10, in dem die Dienste Fernprozeduraufrufe enthalten und die Server-Einträge Schnittstellennamen für die Fernprozeduraufrufe enthalten.

13. Namensdiensteanbieter nach Anspruch 10, in dem die Einträge jedes der Konfigurationsprofile eine nach Priorität geordnete Suchliste bilden.

14. Verfahren zum Betreiben eines verteilten Computersystems, mit den Schritten:

Anschließen einer Anzahl von Clients und einer Anzahl von Diensteanbietern an ein Kommunikationsnetz, wobei jeder Diensteanbieter einen oder mehrere Dienste liefert, die durch einen Dienstnamen identifiziert sind, wobei jeder Client über das Netz unter Verwendung einer eindeutigen Netzadresse Daten an jeden der Diensteanbieter senden und Daten von jedem der Diensteanbieter empfangen kann;

Anbieten eines Namensdienstes als einen Dienst der Diensteanbieter und Speichern von Einträgen für jeden der Diensteanbieter im Namensdienst;

Empfangen von Abfragen von den Clients hinsichtlich der Identität der Diensteanbieter, die durch den Dienstnamen bezeichnet sind, durch

- den Namensdienst und als Antwort auf die Abfragen Zurückschicken von Informationen vom Namensdienst an die Clients, die die Diensteanbieter durch die Netzadresse identifizieren; 5
- wobei der Namensdienst ein Konfigurationsprofil für wenigstens einen Client oder eine Gruppe von Clients bereitstellt, wobei jedes Konfigurationsprofil mehrere Einträge speichert, einschließlich: 10
- optional eines oder mehrerer Server-Einträge, wovon jeder eine Netzadresse eines der Diensteanbieter definiert; 15
- optional einer oder mehrerer Dienstegruppen, wobei jede Dienstegruppe mehrere der Server-Einträge definiert; und
- optional eines oder mehrerer anderer Konfigurationsprofile; 20
- und Speichern von Prioritätsinformationen in jedem der Einträge des Konfigurationsprofils des Namensdiensts, um die Reihenfolge der Verwendung der Netzadressen durch den Client zu bestimmen. 25
15. Verfahren nach Anspruch 14, in dem das Kommunikationsnetz ein lokales Netz und ein Weitverkehrsnetz umfaßt und in dem wenigstens einige der Dienste im lokalen Netz und/oder im Weitverkehrsnetz dupliziert werden. 30
16. Verfahren nach Anspruch 14, in dem die Dienste Fernprozeduraufrufe enthalten und die Server-Einträge Schnittstellennamen für die Fernprozeduraufrufe enthalten. 35
17. Verfahren nach Anspruch 14, in dem die Einträge des Konfigurationsprofils eine nach Priorität geordnete Suchliste bilden. 40
18. Kommunikationsnetz, mit: 45
- a) einem Kommunikationsglied;
- b) mehreren Clients, die an das Kommunikationsglied angeschlossen sind; 50
- c) mehreren Diensteanbietern, die an das Kommunikationsnetz angeschlossen sind;
- d) wobei jeder der Clients und jeder der Diensteanbieter Verarbeitungseinrichtungen enthält, die einen getrennten Befehlsstrom ausführen, und jeder der Clients und jeder der Diensteanbieter eine eindeutige Netzadresse 55
- im Netz besitzt;
- e) wobei jeder der Clients und jeder der Diensteanbieter eine Einrichtung besitzt zum Tätigen von Fernprozeduraufrufen an die Diensteanbieter über das Kommunikationsglied bei Ausführung des Befehlsstroms;
- f) wobei jeder der Clients und jeder der Diensteanbieter eine Einrichtung besitzt zum Senden und Empfangen von Nachrichten an andere bzw. von anderen Clients und Diensteanbietern über das Kommunikationsglied unter Verwendung der Netzadressen, wobei die Nachrichten beim Tätigen der Fernprozeduraufrufe verwendet werden;
- g) wobei einer der Diensteanbieter ein Namensdienst ist, der in der Weise arbeitet, daß er eine die Netzadresse eines besonderen der Diensteanbieter enthaltende Nachricht als Antwort auf eine Nachricht von einem der Clients, die den Namen des Dienstes angibt, zurückschickt;
- h) und einer Einrichtung, die im Namensdienst enthalten ist, um ein Konfigurationsprofil für wenigstens einen der Clients oder eine Gruppe von Clients bereitzustellen, wobei jedes Konfigurationsprofil eine nach Priorität geordnete Suchliste enthält, die Dienstnamen auf Diensteanbieter abbildet, wobei jedes Konfigurationsprofil enthält:
- optional einen oder mehrere Server-Einträge, wovon jeder eine Netzadresse eines Diensteanbieters für einen gegebenen Dienstnamen definiert;
- optional eine oder mehrere mit Namen versehene Dienstegruppen, wovon jede mehrere der Server-Einträge definiert; und
- eines oder mehrere andere Konfigurationsprofile;
- i) einer Einrichtung zum Zuordnen der Konfigurationsprofile zu einem gegebenen Client;
- j) und einer Einrichtung im Namensdienst, um bei Empfang von Abfragen von dem gegebenen Client durch den Namensdienst die Netzadressen zum gegebenen Client zurückzuschicken, um von dem Client entsprechend der Prioritätsreihenfolge der Liste verwendet zu werden.
19. Netz nach Anspruch 18, in dem der Diensteanbieter

- im Diensteeintrag als Fernprozeduraufruf-Schnittstellenspezifikation gekennzeichnet ist.
20. Netz nach Anspruch 18, in dem das Kommunikationsnetz ein lokales Netz und ein Weitverkehrsnetz umfaßt und in dem wenigstens einige der Dienste im lokalen Netz und/oder im Weitverkehrsnetz dupliziert werden.
21. Netz nach Anspruch 18, in dem die Dienstenamen alphabetische Namen und numerische Namen enthalten.

Revendications

1. Procédé pour faire fonctionner un système de traitement distribué, comprenant les étapes qui consistent à :

connecter un certain nombre de clients et un certain nombre de moyens fournisseurs de services à un réseau de communication, chaque moyen fournisseur de services fournissant un service ayant un nom de service, où chaque client est capable d'envoyer des données à et de recevoir des données de chacun desdits moyens fournisseurs de services par l'intermédiaire dudit réseau ;

maintenir un service de désignation comme l'un desdits moyens fournisseurs de services, ledit service de désignation contenant au moins un profil de configuration associé à au moins l'un desdits clients, chaque profil de configuration incluant une liste de recherche de priorités ordonnées faisant correspondre des noms de service à des moyens fournisseurs de services ;

définir pour au moins l'un desdits clients l'un sélectionné desdits profils de configuration pour une utilisation par ledit client ;

faire une demande par l'un desdits clients audit service de désignation en utilisant l'un desdits noms de service, et renvoyer audit client des informations dudit sélectionné desdits profils de configuration pour sélectionner dans ladite liste de recherche l'un desdits moyens fournisseurs de services et fournir ainsi audit client une adresse de réseau dudit moyen fournisseur de services sélectionné ;

faire une demande par ledit client audit moyen fournisseur de services sélectionné pour ledit service.

2. Procédé selon la revendication 1, dans lequel lesdits services incluent des appels de procédure à distance et lesdits moyens fournisseurs de services pour des appels de procédure à distance sont re-

présentés comme des spécifications d'interface dans ledit profil de configuration.

3. Procédé selon la revendication 1, dans lequel chacun desdits profils de configuration comprend une pluralité d'entrées incluant :

en option, une ou plusieurs entrées dans un serveur définissant chacun une adresse de réseau de l'un desdits moyens fournisseurs de services ;

en option, un ou plusieurs groupes de services, chaque groupe de services définissant une pluralité desdites entrées dans un serveur ; et,

en option, un ou plusieurs autres profils de configuration.

4. Procédé selon la revendication 1, dans lequel lesdits noms de service comprennent des noms alphabétiques et des noms numériques.

5. Procédé selon la revendication 1, dans lequel ledit réseau de communication comprend un réseau local et un réseau longue distance, et dans lequel au moins certains desdits services sont reproduits dans ledit réseau local et/ou dans ledit réseau longue distance.

6. Système de traitement distribué, comprenant :

un réseau de communication comportant un certain nombre de clients et un certain nombre de moyens fournisseurs de services connectés à ceux-ci, dans lequel chaque client est capable d'envoyer des données à et de recevoir des données de chacun desdits moyens fournisseurs de services par l'intermédiaire dudit réseau en utilisant une adresse de réseau unique ;

l'un desdits moyens fournisseurs de services fournissant un service de désignation, ledit service de désignation étant un service de base de données contenant des entrées pour chacun desdits moyens fournisseurs de services et recevant des interrogations desdits clients sur l'identité desdits moyens fournisseurs de services indiqués par un nom de service et renvoyant des informations identifiant lesdits moyens fournisseurs de services par une adresse de réseau ;

ledit service de désignation incluant dans ladite base de données un profil de configuration pour au moins un client ou un groupe de clients, chaque profil de configuration comprenant une pluralité d'entrées incluant :

en option, une ou plusieurs entrées dans un serveur définissant chacune une adres-

se de réseau de l'un desdits moyens fournisseurs de services ;
 en option, un ou plusieurs groupes de services, chaque groupe de services définissant une pluralité desdites entrées dans un serveur ; et,
 en option, un ou plusieurs autres profils de configuration ;

chacune desdites entrées de chaque dit profil de configuration dudit service de désignation contenant des informations de priorité pour déterminer l'ordre d'utilisation desdites adresses de réseau par ledit client après réception desdites interrogations par ledit service de désignation ;
 et des moyens pour identifier pour ledit client un ou plusieurs sélectionnés desdits profils de configuration pour qu'il s'y réfère en faisant lesdites interrogations.

7. Système selon la revendication 6, dans lequel ledit réseau de communication comprend un réseau local et un réseau longue distance, et dans lequel au moins certains desdits services sont reproduits dans ledit réseau local et/ou dans ledit réseau longue distance.

8. Système selon la revendication 6, dans lequel lesdits services comprennent des appels de procédure à distance, et lesdites entrées dans un serveur comprennent des noms d'interface pour lesdits appels de procédure à distance.

9. Système selon la revendication 6, dans lequel lesdites entrées de chaque dit profil de configuration constituent une liste de recherche de priorités ordonnées.

10. Moyen fournisseur de service de désignation pour système de traitement distribué, le système comprenant un réseau de communication comportant un certain nombre de clients et un certain nombre de moyens fournisseurs de services connectés à ceux-ci, chacun desdits moyens fournisseurs de services fournissant un service identifié par un nom de service, ledit moyen fournisseur de service de désignation comprenant :

un service de base de données contenant des entrées pour chacun d'une pluralité de services fournis par lesdits moyens fournisseurs de services ;
 des moyens pour recevoir des interrogations desdits clients sur l'identité desdits moyens fournisseurs de services indiqués par un nom de service et renvoyant des informations identifiant lesdits moyens fournisseurs de services

par une adresse de réseau ;
 lesdits moyens incluant un profil de configuration inclus dans ladite base de données pour l'un desdits clients ou pour un groupe desdits clients, chaque dit profil de configuration comprenant une pluralité d'entrées incluant :

en option, une ou plusieurs entrées dans un serveur définissant chacune une adresse de réseau de l'un desdits moyens fournisseurs de services ;
 en option, un ou plusieurs groupes de services, chaque groupe de services définissant une pluralité desdites entrées dans un serveur ; et,
 en option, un ou plusieurs autres profils de configuration ;

chacune desdites entrées de chaque dit profil de configuration dudit service de désignation contenant des informations de priorité pour déterminer l'ordre d'utilisation desdites adresses de réseau par ledit client ;
 et des moyens pour identifier pour ledit client un ou plusieurs sélectionnés desdits profils de configuration pour s'y référer en faisant lesdites interrogations.

11. Moyen fournisseur de service de désignation selon la revendication 10, en combinaison avec un réseau de communication incluant un réseau local et un réseau longue distance, et dans lequel au moins certains desdits services sont reproduits dans ledit réseau local et/ou dans ledit réseau longue distance.

12. Moyen fournisseur de service de désignation selon la revendication 10, dans lequel lesdits services comprennent des appels de procédure à distance, et lesdites entrées dans un serveur comprennent des noms d'interface pour lesdits appels de procédure à distance.

13. Moyens fournisseur de service de désignation selon la revendication 10, dans lequel lesdites entrées de chaque dit profil de configuration constituent une liste de recherche de priorités ordonnées.

14. Procédé pour faire fonctionner un système de traitement distribué, comprenant les étapes qui constituent à :

connecter un certain nombre de clients et un certain nombre de moyens fournisseurs de services à un réseau de communication, chaque moyen fournisseur de services fournissant un ou plusieurs services identifiés par un nom de service, dans lequel chaque client est capable

d'envoyer des données à et de recevoir des données de chacun desdits moyens fournisseurs de services par l'intermédiaire dudit réseau en utilisant une adresse de réseau unique ;

fournir un service de désignation comme l'un desdits moyens fournisseurs de services, et enregistrer dans ledit service de désignation des entrées pour chacun desdits moyens fournisseurs de services ;

recevoir par ledit service de désignation des interrogations desdits clients sur l'identité desdits moyens fournisseurs de services indiqués par un nom de service et renvoyer des informations dudit service de désignation auxdits clients identifiant lesdits moyens fournisseurs de services par une adresse de réseau en réponse auxdites interrogations ;

ledit service de désignation maintenant un profil de configuration pour au moins un client ou un groupe de clients, chaque profil de configuration enregistrant une pluralité d'entrées incluant :

en option, une ou plusieurs entrées dans un serveur définissant chacune une adresse de réseau de l'un desdits moyens fournisseurs de services ;

en option, un ou plusieurs groupes de services, chaque groupe de services définissant une pluralité desdites entrées dans un serveur ; et,

en option, un ou plusieurs autres profils de configuration ;

et enregistrer dans chacune desdites entrées dudit profil de configuration dudit service de désignation des informations de priorité pour déterminer l'ordre d'utilisation desdites adresses de réseau par ledit client.

15. Procédé selon la revendication 14, dans lequel ledit réseau de communication comprend un réseau local et un réseau longue distance, et dans lequel au moins certains desdits services sont reproduits dans ledit réseau local et/ou dans ledit réseau longue distance.

16. Procédé selon la revendication 14, dans lequel lesdits services comprennent des appels de procédure à distance, et lesdites entrées dans un serveur comprennent des noms d'interface pour lesdits appels de procédure à distance.

17. Procédé selon la revendication 14, dans lequel lesdites entrées dudit profil de configuration constituent une liste de recherche de priorités ordonnées.

18. Réseau de communication comprenant :

a) une liaison ;

b) une pluralité de clients couplés à ladite liaison ;

c) une pluralité de moyens fournisseurs de services couplés audit réseau de communication ;

d) chacun desdits clients et chacun desdits moyens fournisseurs de services incluant des moyens de traitement exécutant une suite d'instructions séparée, et chacun desdits clients et chacun desdits moyens fournisseurs de services ayant une adresse de réseau unique dans ledit réseau ;

e) chacun desdits clients et chacun desdits moyens fournisseurs de services comportant des moyens pour faire des appels de procédure à distance auxdits moyens fournisseurs de services par ladite liaison, quand ils exécutent ladite suite d'instructions ;

f) chacun desdits clients et chacun desdits moyens fournisseurs de services comportant des moyens pour émettre et recevoir des messages pour d'autres clients et moyens fournisseurs de services par ladite liaison en utilisant lesdites adresses de réseau, lesdits messages étant mis en oeuvre en faisant lesdits appels de procédure à distance ;

g) l'un desdits moyens fournisseurs de services étant un service de désignation fonctionnant pour renvoyer un message contenant l'adresse de réseau de l'un particulier desdits moyens fournisseurs de services en réponse à un message provenant de l'un desdits clients donnant le nom dudit service ;

h) et des moyens inclus dans ledit service de désignation pour maintenir un profil de configuration pour au moins l'un desdits clients ou un groupe de clients, chaque profil de configuration comprenant une liste de recherche de priorités ordonnées faisant correspondre des noms de service dans les moyens fournisseurs de services, chaque profil de configuration incluant :

en option, une ou plusieurs entrées dans un serveur définissant chacune une adresse de réseau d'un moyen fournisseur de services pour un nom de service donné ;

en option, un ou plusieurs groupes de services désignés définissant chacun une pluralité desdites entrées dans un serveur ; et, en option, un ou plusieurs autres dits profils de configuration ;

i) des moyens pour associer l'un desdits profils de configuration à un client donné ;

j) et des moyens dans ledit service de désigna-

tion pour renvoyer lesdites adresses de réseau audit client donné à la réception par ledit service de désignation d'interrogations dudit client donné, pour une utilisation par ledit client selon l'ordre de priorité de ladite liste.

5

19. Réseau selon la revendication 18, dans lequel ledit moyen fournisseur de services est identifié dans ladite entrée dans un serveur comme spécification d'interface d'appel de procédure à distance.

10

20. Réseau selon la revendication 18, dans lequel ledit réseau de communication comprend un réseau local et un réseau longue distance, et dans lequel au moins certains desdits services sont reproduits dans ledit réseau local et/ou dans ledit réseau longue distance.

15

21. Réseau selon la revendication 18, dans lequel lesdits noms de service comprennent des noms alphabétiques et des noms numériques.

20

25

30

35

40

45

50

55




FIG 1




FIG 2

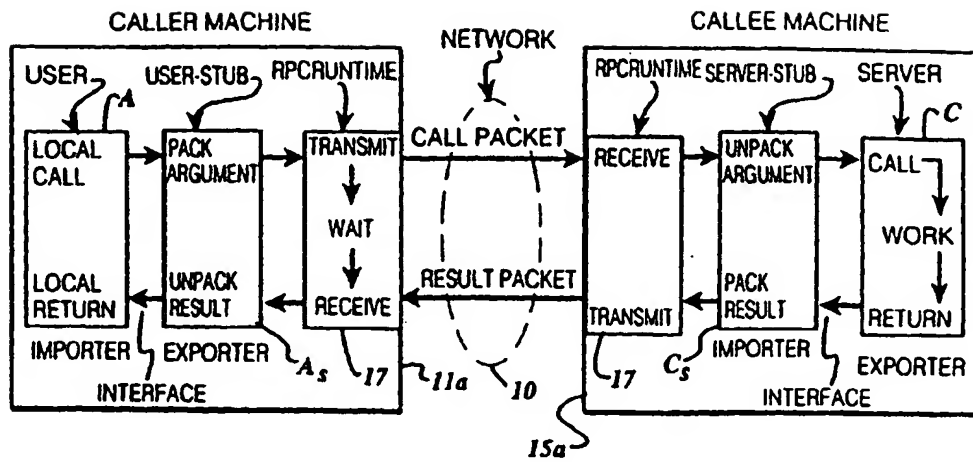


FIG 3

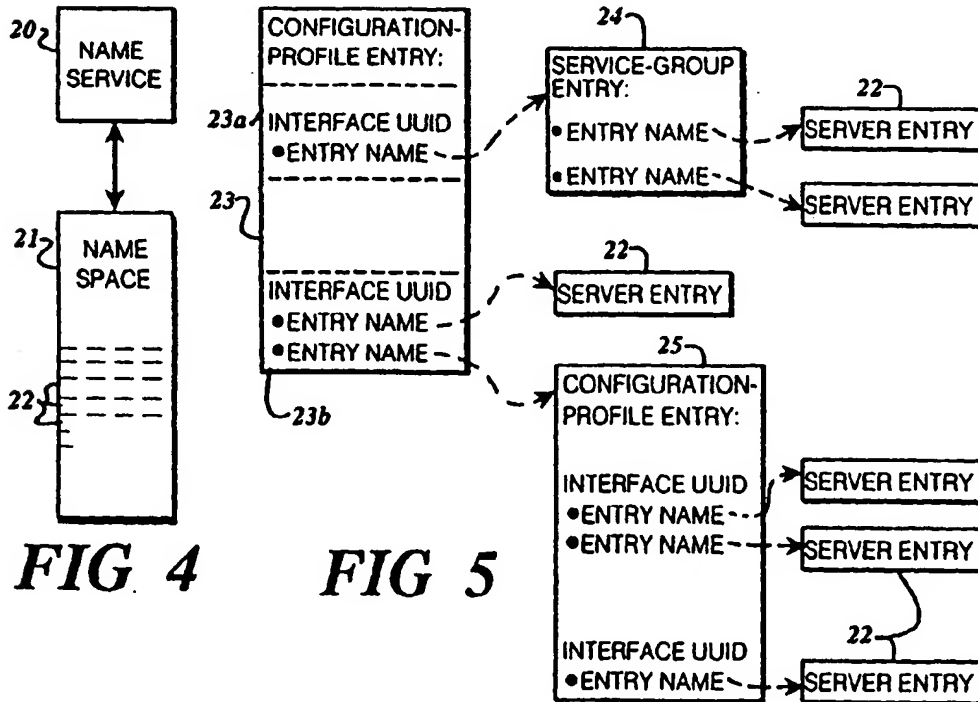


FIG 4

FIG 5

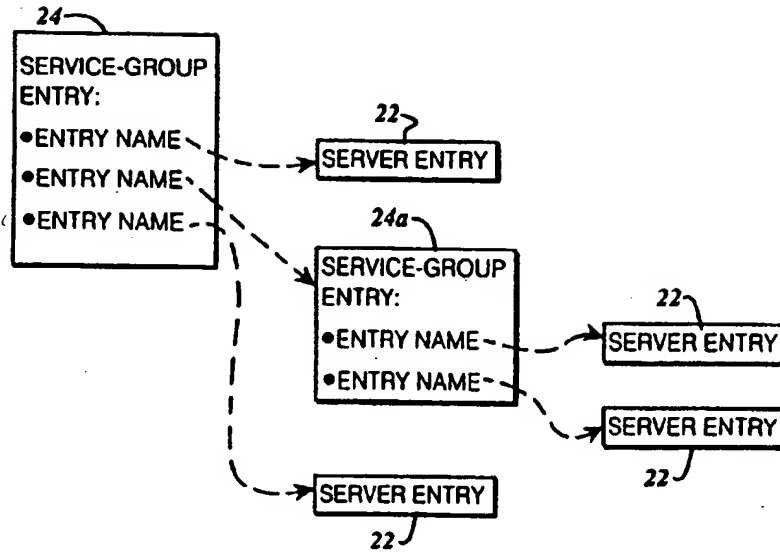


FIG 6




FIG 7




FIG 7a